tor (the ICTZ) at the latitude of the Cariaco Basin. This is also the latitude of the Panama Isthmus, where maximum transfer of water vapor to the Pacific Ocean could occur, thus increasing the mean salinity of the Atlantic warm pool. Warm D-O events would then be reinforced by the increased salt content in the source water for the North Atlantic thermohaline convection.

This is an interesting hypothesis, but this feedback effect would probably be negligible over the short warming phase (see the figure). The temporal relation between low-latitude forcing and the evolution of the northern ice sheets, wind field, and thermohaline circulation is therefore still not known at sufficient resolution (a few decades) to establish the D-O mechanism. We need to use better paleoclimatic proxies, especially to quantify sea surface temperature, and more high sedimentation rate cores, in particular in the Gulf of Mexico. Hydrological changes in the Gulf Stream source waters must also be characterized better if we want to understand the role of the thermohaline circulation in rapid climate change.

References and Notes

1. K. A. Hughen, J. R. Southon, S. J. Lehman, J. T. Overpeck, *Science* **290**, 1951 (2000).
2. L. C. Peterson, G. H. Haug, K. A. Hughen, U. Röhl, *Science* **290**, 1947 (2000).
3. G. Bond *et al.*, *Nature* **365**, 143 (1993).
4. E. Cortijo *et al.*, *Earth Planet. Sci. Lett.* **146**, 29 (1997).
5. L. Vidal *et al.*, *Earth Planet. Sci. Lett.* **146**, 13 (1997).
6. T. L. Rassmussen, T. C. E. Van Weering, L. Labeyrie, *Quat. Sci. Rev.* **16**, 71 (1997).
7. C. Rühlemann *et al.*, *Nature* **402**, 511 (1999).
8. See http://www.images.cnrs-gif.fr
9. K. A. Hughen, J. T. Overpeck, L. C. Peterson, S. Trumbore, *Nature* **380**, 51 (1996).
10. P. M. Grootes, M. Stuiver, *J. Geophys. Res.* **102**, 26455 (1997).
11. P. M. Grootes *et al.*, *Nature* **366**, 552 (1993).
12. I. L. Hendy, J. P. Kennett, *Geology* **27**, 291 (1999).
13. N. J. Shackleton, M. A. Hall, E. Vincent, *Paleoceanography*, in press.
14. S. van Krevelt *et al.*, *Paleoceanography* **15**, 425 (2000).

15. Institut Français de Recherche et Technologie Polaire is acknowledged for its large support of the IMAGES Marion Dufresne (MD) cruises.

PERSPECTIVES: OPTICS

The Internet of Tomorrow

Steve Joiner

The digital economy is generating an insatiable appetite for bandwidth. Internet traffic now hits daily peaks of 0.47 terabits per second and is expected to increase by more than an order of magnitude to 5.64 terabits per second by the end of 2001 (*1*). The data moved around the globe are getting richer and more dynamic with each passing day, and an increasingly impatient public expects faster and faster responses. This endless demand is creating stress points that threaten to make the Web's infrastructure collapse under the weight of its own success.

The problem does not lie with data transfer from city to city or even across continents or oceans. Because of advances in dense wave-division multiplexing technology, which enables the creation of multiple communication paths on one fiber through the use of multiple wavelengths of light in the fiber, fiber communication capacity has doubled every nine months. The increased capacity of the fiber enables much more efficient use of installed fiber, thus lowering operating costs for long-haul optical communication.

The real bottleneck today is the access network that extends from the desktop to the service provider's Internet point of presence (see the figure). This "short reach" or "last mile" still uses the installed base of traditional copper connections or the copper used in cable TV systems. The new high-speed Internet connections based on cable modems or digital subscriber line (DSL) both rely on fiber connections terminating relatively close to the end user. The optical connections used in these cases are becoming more economical as volumes increase. Thus, optics is beginning to solve today's "last-mile" bandwidth bottleneck.

In the near future, the bottlenecks choking the Internet will be even closer to home: the very short-reach interfaces that connect buildings, floors, racks, storage subsystems, internal computer components, and even the onboard elements of integrated circuits. If the Internet is to continue its explosive growth, optical elements must be scaled down and applied at these levels. One of the most promising technologies is a relative newcomer, the vertical cavity surface emitting laser (VCSEL).

Traditional lasers are edge-emitting, whereas VCSELs emit light from their top surface in a cylindrical beam. This has the advantage that the emitted light requires no corrections for asymmetry or astigmatism and can be shaped into a ring for various applications. VCSELs are also very efficient, requiring just a fraction of the power used by edge-emitting lasers, thus simplifying the electronic circuitry that drives them and reducing radio frequency interference (RFI) emissions. And VSCELs are smaller than conventional lasers: The entire VCSEL can be shrunk until it is only slightly larger than the beam size.

Surface emission also allows for easier mounting and packaging. VCSELs can be attached flat onto the package substrate by the same type of equipment used to assemble integrated circuits and can be tested directly on the wafers. Separating the lasers from the wafers and packaging them into arrays becomes trivial, the yield per wafer is a much higher, and the packing density increases by an order of magnitude—a critical capability for constructing high-speed optical interconnects for computer systems and networks. All these factors reduce development and manufacturing costs and move the benefits of optics closer to the consumers.

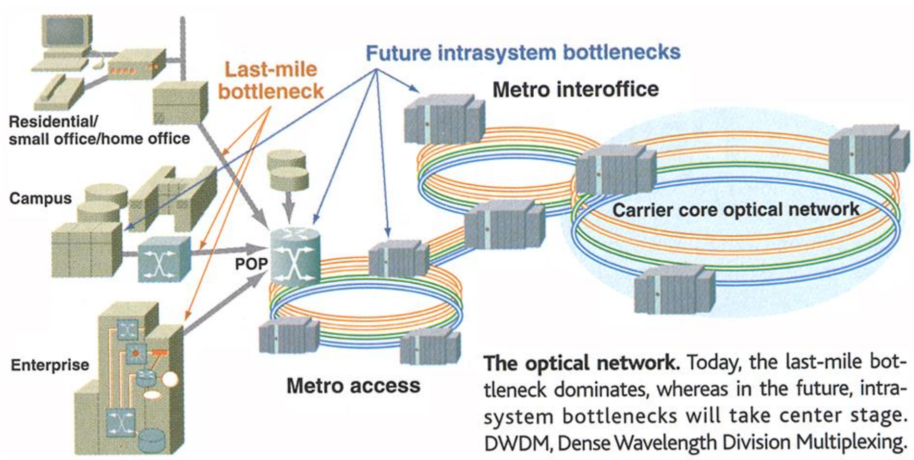

The optical network. Today, the last-mile bottleneck dominates, whereas in the future, intrasystem bottlenecks will take center stage. DWDM, Dense Wavelength Division Multiplexing.

As the data traffic from the desktop is routed through the network and aggregated with other traffic in the backbone of the Internet, the traffic inside the routing switches will increase to the point that electrical connections will not work even over distances covering a single piece of networking equipment. Manufacturers are turning

The author is with Agilent Laboratories, 3500 Deer Creek Road, MS 26M, Palo Alto, CA 94304, USA. E-mail: steve_joiner@agilent.com

to VCSEL-based optical interconnections inside these large telecommunication switches to move data between PC boards.

When VCSELs can be commercially made to emit in the 1300- to 1550-nm range, they will also be useful with the single-mode fiber installed in the "last-mile" bottleneck mentioned earlier. Thus, the simplification of optical links enabled by the VCSEL will be applied to the short-reach single-mode connections, further reducing costs as the optical network migrates closer to the home. Several companies have presented data on such long-wavelength VCSEL prototypes.

Speeding up electronic routers with optical backplanes is, however, only a temporary solution. Since the invention of the semiconductor chip, the number of circuits on a chip has doubled every 18 months, following Moore's law. This results in denser and denser printed circuit boards, and the communications channels that move the signals between components are becoming overloaded. Integrated circuits will soon generate so much input/output (I/O) traffic that they won't be able to communicate electronically. VCSELs can be applied at this level in a "photonics on CMOS" architecture that will enable Internet router throughput to continue to increase (*2*, *3*).

As these developments make access networks faster and faster, the I/O bus of the desktop computer becomes the Internet's choke point. Ever more powerful CPUs cannot move the content off the network and into the display subsystem fast enough because of this internal constraint. You can see a microcosm of this problem when you send a document to a locally attached printer and try to work on another document while it prints. There is plenty of processor power available, but your computer responds sluggishly because there is no room on the bus. Network servers have the same problem, and the computer industry is addressing it with Infiniband, a switched fabric I/O architecture that gets rid of the old shared-bus architecture (*4*). Think of it as an Internet inside a single computer. If VCSELs become cheap enough, these intracomputer networks can become optical, thus solving the I/O bottleneck on a single PC board.

Before this can happen, though, the optical industry must overcome some key challenges that are encountered when optical interfaces are placed directly on a microprocessor. First, there is the issue of heat. Integrated circuits run at a much higher temperature than VCSELs are designed to withstand, creating a materials problem. Second, there is a form-factor problem. The electrical-to-optical converter needs to be as small as two wire bond pads on an integrated circuit, or about 0.03 mm^2. Today, the equivalent function requires a 12 mm by 50 mm package that fills up a 600 mm^2 area on a mounting surface. The optical technology thus needs to shrink by four orders of magnitude. Third, the manufacturing yields of optical transceivers need to improve by about two orders of magnitude.

Optics has solved the communication bottlenecks for undersea and long-haul applications. Reduced cost of short-reach optical solutions is now bringing optics closer to the home to solve the last-mile bottleneck of today. VCSELs will be the enabling technology to solve the emerging bottlenecks inside computers and telecommunications equipment. Eventually, VCSELs will be needed to communicate across individual PC boards. As the demand for bandwidth increases and optical link technology evolves to produce lower cost solutions, optics is set to solve each bottleneck in our internet infrastructure.

References and Notes

1. According to telecommunications industry research from Ryan Hankin Kent Inc. (www.rhk.com), 2000.
2. R. C. Walker, paper presented at 10th Annual Workshop on Interconnections within High-Speed Digital Systems, Santa Fe, NM, 9 May 1999.
3. W. Ishak, S. Joiner, paper presented at Workshop on DARPA Initiative on Information Technology, Santa Fe, NM, 12 May 1999.
4. See www.infinibandta.org

PERSPECTIVES: NEUROSCIENCE

Noise Makes Sense in Neuronal Computing

Maxim Volgushev and Ulf T. Eysel

Everyone agrees that a noisy telephone line does not improve the transmission of information. Yet, under certain circumstances, added noise can actually enhance rather than diminish the detection of weak signals. Neurons in the primary visual cortex, the part of the brain responsible for processing incoming visual signals, seem to have taken advantage of this effect, as Anderson *et al.* (*1*) show on page 1968 of this issue. They present new evidence that synaptic noise in primary visual cortex neurons enables these cells to maintain their encoding of the orientation of lines in an image at the same quality, regardless of whether the image is seen at high or low contrast (a phenomenon called contrast invariance of orientation tuning).

Nerve cells use action potentials (also called "discharges" or "spikes") to encode information that they transmit to other neurons. In the visual system, this information contains many different aspects of a viewed image including color, contrast, and orientation of lines and borders of the image. Generation of action potentials by a primary visual cortex neuron critically depends on parts of the visual image being orientated in the preferred range for that neuron. Visual stimulus orientations in the range preferred by that cell evoke action potentials, whereas other orientations do not (*2*). But these other orientations do elicit responses in other neurons. At the same time, the rate at which action potentials are produced increases in all neurons of the primary visual cortex when the contrast of the image is increased. Although the response strength does change as the stimulus contrast is varied, the tuning of a cell (that is, the dependence of the action potential frequency on the orientation of the stimulus) does not—a property called contrast invariance (*3*). This invariance allows independent encoding of the two different stimulus features, namely, contrast and orientation, by the same neuron.

Anderson *et al.* (*1*) report that contrast invariance of orientation tuning holds true not only for the action potentials generated, but also for the membrane potential responses, that is, at low contrast the signal is weak but the tuning width for stimulus orientation is preserved (see the figure). This result is in apparent conflict with a nerve cell's dependence on its membrane potential reaching a threshold before an action potential is generated. Classically, a neuron is considered to be an integrate-and-fire unit, with a simple relation between the depolarization of its membrane (due to a net influx of positively charged ions into the cell) and generation of the action potential. When the membrane depolarization reaches a certain threshold, action potentials are generated, and with a further increase in depolarization the rate of action potential generation increases (*4*). The necessity for the membrane potential to reach a threshold before being transformed into a train of action potentials accounts for sharpening of orientation tuning (*5*) through the "Iceberg effect" (see the figure, left). Like the tip of an iceberg, only part of the modulated membrane potential is visible above the threshold (red line in the figure); accordingly, action potential responses are

The authors are in the Department of Neurophysiology, Faculty of Medicine, Ruhr-University Bochum, D-44780 Bochum, Germany. E-mail: eysel@neurop.ruhr-uni-bochum.de